

BBC

Research Department Report

October 1987


DIGITAL VIDEO BIT-RATE REDUCTION: DPCM for sub-Nyquist PAL

J.H. Stott, M.A.

DIGITAL VIDEO BIT-RATE REDUCTION: DPCM FOR SUB-NYQUIST PAL**J.H. Stott, M.A.****Summary**

Differential pulse-code modulation (DPCM) is a technique for bit-rate reduction. This Report discusses the principles of DPCM and its application to television. An investigation of the use of DPCM with sub-Nyquist sampled composite PAL video signals is then described. The usefulness of computer optimisation techniques is established and practical development work described. It is concluded that three-dimensional predictors, of a complexity not previously envisaged to be worthwhile, offer significant improvements. Two systems meeting practical requirements for 34 and 68 Mbit/s transmission 'packages' were developed and subjectively assessed.

Issued under the Authority of

**Research Department, Engineering Division,
BRITISH BROADCASTING CORPORATION**

Head of Research Department

This Report may not be reproduced in any form
without the written permission of the
British Broadcasting Corporation

It uses SI units in accordance with B.S. document
PD 5686

UNIVERSAL DECIMAL CLASSIFICATION

It is intended that the UDC reference (see Summary page, top right) will, from 1988 onwards, be omitted from BBC Research Reports. The classification is not generally used elsewhere within the BBC. Should, however, the lack of this reference be of particular concern to you would you kindly write to:

The Research Executive,
BBC Research Department, Kingswood Warren,
Tadworth, Surrey, KT20 6NP, England.

DIGITAL VIDEO BIT-RATE REDUCTION: DPCM FOR SUB-NYQUIST PAL

J.H. Stott, M.A.

1. Introduction	1
2. What is DPCM?	1
2.1 Properties of the signal source	1
2.2 Variable-length coding (v.l.c.) — a reversible technique	2
2.3 DPCM — a non-reversible technique	3
2.4 A practical DPCM arrangement	4
2.5 Non-causal prediction	5
2.6 Adaptation of predictors and/or quantisers	5
3. Application of DPCM to Digital Video — Background	6
3.1 DPCM applied to monochrome signals	6
3.2 DPCM applied to colour signals	7
3.3 Scope for improvement	7
4. Development Tools	7
4.1 The approach to development	7
4.2 Software tools	8
4.3 Hardware tools	9
5. Development Work	9
5.1 Computed prediction gains for 'optimum' predictors	9
5.2 Prediction optimisation under constraints	10
5.3 Quantisers	12
5.4 Resetting, rounding error, 'smoke' and the use of rounding error feedback	13
5.5 Systems selected for practical applications	16
6. Further Work	17
7. Subjective Tests	17
7.1 The scope of the tests	17
7.2 Subjective test method	19
7.3 Experiment 1: moving pictures, colour viewing, PCM reference	20
7.4 Experiment 2: moving pictures, colour viewing, analogue reference	20
7.5 Experiment 3: moving pictures, monochrome viewing, PCM reference	20
7.6 Experiment 4: still pictures, colour viewing, analogue reference	21
7.7 Experiment 5: still pictures, monochrome viewing, analogue reference	21
7.8 Experiment 6: effect of transmission errors	21
8. Conclusions	21
9. Acknowledgments	22
10. References	22
Appendix	24

© BBC 2006. All rights reserved. Except as provided below, no part of this document may be reproduced in any material form (including photocopying or storing it in any medium by electronic means) without the prior written permission of BBC Research & Development except in accordance with the provisions of the (UK) Copyright, Designs and Patents Act 1988.

The BBC grants permission to individuals and organisations to make copies of the entire document (including this copyright notice) for their own internal use. No copies of this document may be published, distributed or made available to third parties whether by paper, electronic or other means without the BBC's prior written permission. Where necessary, third parties should be directed to the relevant page on BBC's website at <http://www.bbc.co.uk/rd/pubs/> for a copy of this document.

DIGITAL VIDEO BIT-RATE REDUCTION: DPCM FOR SUB-NYQUIST PAL

J.H. Stott, M.A.

1. INTRODUCTION

The advantages of digital transmission are well known, namely that digitally-coded signals are inherently highly resistant to the effects of noise and distortion encountered in transmission systems and can be regenerated for long-distance transmission. Accordingly the quality of television signals transmitted in digital form is usually determined by the coding methods used rather than the link itself. Digital transmission of video signals is therefore attractive, the more so as other factors come into play. Digital networks, in which digital television can take its place alongside telephony and data without special treatment, are being established in the UK by the telecommunications authorities. Indeed, the provision of analogue video links is becoming the special, and thus expensive, case. At the same time, as increasing use is being made of digital processes for programme origination, digital transmission between studios or to broadcast transmitters would form a logical step in the process of evolution towards all-digital operation.

However, straightforward digital coding of video signals generates a signal having a high bit rate (over 100 Mbit/s) which, if transmitted, is rather wasteful of channel capacity. Possible ways of reducing the bit rate required for video signals have therefore received much attention. Previous work in BBC Research Department produced an experimental system which has been used in several field trials of digital transmission^{1,2,3,4}. One video signal and six high-quality sound channels were coded⁵ and multiplexed⁶ into a 60 Mbit/s 'package'; two such packages could be multiplexed into 120 Mbit/s.

To be useful a package must provide a suitable combination of video and other signals within a bit rate at which access to the digital network is permitted. The figures of 120 and 60 Mbit/s mentioned above were appropriate for the experimental digital cable and satellite systems of the time, but since then a hierarchy of bit rates has been agreed within Europe. This uses the following bit rates: 2.048, 8.448, 34.368 and 139.264 Mbit/s. The 139.264 Mbit/s bit rate (commonly called '140 Mbit/s') may not be universally available in Europe, while in the UK access may also be arranged at 68.736 Mbit/s ('68 Mbit/s') which is exactly twice 34.368 Mbit/s ('34 Mbit/s').

An obvious target for the next stage of bit-rate reduction studies was 34 Mbit/s. It appeared to be attainable using techniques which were an extension of

those gained from previous BBC experience. The knowledge gained would also be useful for other, higher, bit rates if they proved to be economically more attractive, bearing in mind the balance between circuit and equipment costs. Conversely, to reduce the bit rate of a broadcast-quality video signal to fit within 8 Mbit/s would undoubtedly require much more advanced techniques.

Three techniques were studied in the investigation⁷, all three being needed in order to meet the 34 Mbit/s target. They were: line- and field-blanking removal⁸, so-called sub-Nyquist sampling of the PAL signal, and differential pulse-code modulation (DPCM). Sub-Nyquist sampling and DPCM were used in the previously-developed 60 Mbit/s package, with promising results. Another Report⁹ describes the further work on sub-Nyquist sampling whereby the impairments it introduces have been minimised.

The purpose of this Report is to describe the further work on DPCM coding whereby the number of bits per sample has been reduced without sacrificing picture quality.

2. WHAT IS DPCM?

DPCM is a bit-rate reduction technique in which the difference between each input sample and its prediction from previous samples is quantised before transmission such that the number of bit/sample is reduced. It exploits the properties of both the *signal source* and the *signal destination* (usually a human observer). DPCM bears some similarity to certain forms of variable-length coding, and the latter therefore will also be referred to in this section.

2.1 Properties of the signal source

Predictive coding (of which DPCM is an example) exploits the redundancy present in the output of many types of sources whereby the sampled signal values are not random numbers, but are related in some way. Because of this relationship we can make a prediction of the value of a sample from knowledge of *previous** samples.

Consider the properties of a signal called the *prediction error*, formed by subtracting a prediction from the actual value of each sample, as illustrated in Fig. 1.

* This is a causal, extrapolative prediction. Interpolative coding using non-causal prediction is discussed in Section 2.5.




Fig. 1 - Forming a difference signal.

If the predictor were perfect the prediction error would always be zero. This will not be achievable in practice unless the input signal has no information content. It may however be represented by a probability distribution function (p.d.f.) which is zero except at the origin, Fig. 2(a). A practical predictor tailored to a practical signal having some redundancy will give a peaked p.d.f. of the form sketched in Fig. 2(b). A random signal, or a predictor which just guesses, will cause a uniform p.d.f., Fig. 2(c). The original sample values are assumed to take $(R+1)$ different values from 0 to R units (or $-R/2$ to $+R/2$ if the signal is bipolar).




Fig. 2 - Probability-density functions of prediction-error signals.

What does Fig. 2(b) mean? It shows that the prediction error may take $(2R+1)$ different values (i.e. more than for the original signal) but the smaller ones occur more frequently than the large ones. This property can be exploited in two ways for bit-rate reduction, although at first sight it looks as though matters have been made worse: to describe $(2R+1)$ values exactly and directly needs about 1 bit/sample more than to describe the $(R+1)$ values of the source. The two methods are variable-length coding and DPCM; they will be discussed in the following sections.

What form might the prediction take? In the most general sense the predictor is, as we have noted, a device which uses knowledge, of whatever kind, of past sample values to estimate the value of the next.

In the limit this could be an extremely complicated process taking into account all the observed past history of the signal. For example, in television the first frames of commonly repeated sequences — e.g. the graphics sequence used to introduce a regular programme such as the news — could be recognised as they occur whereafter the rest of that particular sequence could be predicted perfectly accurately (until that regularly occurring sequence was modified in some way). This is clearly an extreme which we cannot at present hope to exploit. In practice predictors are normally linear predictors: they form a weighted sum of the values of a limited number of recently occurring samples. Non-linear predictors may also be considered; predictors which use different weightings for the samples according to some function of the signal behaviour could be so described. These are normally considered as *adaptive* predictors.

2.2 Variable-length coding (v.l.c.) — a reversible technique

This is a technique which, like DPCM, takes advantage of a peaked distribution like Fig. 2(b). We discuss it briefly here in order to understand how it differs from DPCM.

Consider the original source under discussion. It generates samples each of which takes one of $(R+1)$ values. We require to send information to a remote decoder so that it may reproduce these sample values. If we send a straightforward, fixed-length binary codeword to describe each sample it will have length $\log_2(R+1)$ bits. If there are f_0 sample/s then the bit rate is clearly $f_0 \log_2(R+1)$ bit/s.

Now suppose that we use codewords of variable length, choosing short codewords for values which occur frequently and long codewords for those values which occur rarely. Clearly, to make this possible we must start with a signal having a peaked distribution like Fig. 2(b); a prediction arrangement like Fig. 1 is used to achieve this if the source does not already have this property. Using the v.l.c. technique the mean bit rate is then reduced, the frequent short codewords outweighing the infrequent long ones. The codewords have to be chosen so that it is possible to separate each one from the next unambiguously despite their variable length; various algorithms have been devised to do this.

The limit to the bit-rate reduction attainable may be expressed in terms of the entropy^{10,11,12} of the signal, H : this is the minimum average number of bit/sample which could be used to describe the signal and is, in effect, a measure of the 'peakiness' of the p.d.f. A random signal has a high entropy, while a signal usually taking only a few values has a low

entropy. The minimum bit rate to transmit the signal is Hf_0 bit/s.

The attraction of this system is that it is *reversible*: the decoder can reproduce the original signal exactly without any loss being introduced by the coding. The system has only exploited the properties of the *source* (which must be known); since no impairment is introduced the properties of the *destination* are of no concern.

Practical systems do not always retain this reversibility in full. Because the codewords have variable length, while a transmission link usually requires a steady flow of information, it is necessary to use a buffer store to smooth out the variations and thus transmit a constant bit rate equal to the average coded bit rate. A practical constraint is imposed by the size of this buffer store. The larger this is, the greater the cost and also the greater the delay in transmission of the signal. It is possible for the store to overflow or underflow during a period when the signal protractedly takes values whose codewords are very long or very short. Underflow may be avoided by a slight modification to the code structure, with little loss of efficiency, but overflow is a more serious problem. The course usually adopted is to have 'fallback' modes of operation in which some impairment of the signal is introduced when the buffer is nearing overflow. These fallback modes are often a form of DPCM (since an appropriate predictor is already being used). Subsampling in one or more dimensions is sometimes used as an alternative or in addition.

To sum up: variable-length coding

- (a) exploits source properties, to which it is matched,
- (b) is reversible in principle, although it —
- (c) in practice introduces some impairment occasionally because of practical constraints.
- (d) has a minimum possible bit rate of Hf_0 , where H is the entropy of the signal encoded and f_0 is the sample rate.

2.3 DPCM — a non-reversible technique

DPCM exploits the peaked p.d.f. of Fig. 2(b) in a different way. Instead of using codes of variable length, as described above for v.l.c., it uses codewords of variable *accuracy*. With this type of coding it is accepted that impairment of the signal will be introduced, but its effect is minimised by concentrating the impairment on those signal features which occur rarely and/or which mask the effects of impairment. Higher accuracy is maintained for commonly-occurring parts of the signal.

The variable accuracy is obtained by the use of a non-linear quantiser (n.l.q.). This has decision levels which are closely spaced in the region of the origin but are progressively more coarsely spaced as the signal amplitude increases. This non-linear operation is applied to the prediction-error signal and the resulting quantised difference values are transmitted to the decoder. Suppose that the quantiser has T levels, i.e. only T distinct values of quantised difference can occur. A simple look-up table can translate these values to transmitted codewords having a fixed length of $\log_2 T$ bit/sample. The bit rate has been reduced provided $T < (R+1)$. A typical non-linear quantiser is shown in Fig. 3.

The greatest impairment is introduced when the prediction-error magnitude is large, which, statistically, occurs rarely. By matching the non-linear quantising characteristics to the p.d.f. characteristics of the source-predictor combination it is possible to minimise the impairment as measured by some objective criterion, such as mean-square error. However, when a system introduces impairments, as this must, it is essential to consider the behaviour of the destination receiving the signal. This might be the human eye or ear, or an intelligent machine, e.g. a pattern recognition device. These different destinations will be disturbed by impairments in different ways, finding some impairments more objectionable than others.




Fig. 3 - Typical n.l.q. characteristic.

A DPCM system thus exploits the properties of both the source and the destination. An ideal system would concentrate impairments on those signal features which occur rarely and which are also those features for which the destination is least sensitive. This is the ideal we aim for in DPCM coding and it requires optimisation of the predictor and the

quantiser. The best predictor for this application is not in general the same as would be used for variable-length coding of the same source. For this reason it may not necessarily be profitable to combine the two techniques.

There is no minimum achievable bit rate when using DPCM. We simply trade bit rate for quality: when the n.l.q. has few levels there is more impairment than when it has many levels. The degree of predictor complexity which is considered worthwhile may vary as the bit rate is varied, since a better predictor may offset the ill effects of reducing the number of n.l.q. levels.

2.4 A practical DPCM arrangement

At first sight it appears that a DPCM coder would consist of a circuit of the form of Fig. 1 with a quantiser added on the end. However, when we attempt to devise a decoder for this arrangement we find that it is not possible, as the decoder does not have the original signal on which to base its prediction. A different structure is necessary in which the prediction is based on past *decoded* values, and this is illustrated in Fig. 4, which also serves to define frequently-used terms.

Apart from the transmission decoder, which is a complementary look-up table to the transmission coder, the decoder forms a subset of the coder circuitry. The coder includes a decoder circuit which forms a locally-decoded signal which is used as the input to the predictor. This feature is useful for development work since coding quality may be

assessed with only a coder, using this locally-decoded output. One word of caution must be sounded, however. The locally-decoded and remotely-decoded outputs are the same only if:

- (a) the initial conditions are the same, and
- (b) no errors occur during transmission of the coded quantised difference signal.

Condition (a) may be satisfied in two ways. One is to reset the contents of the predictors in the coder and decoder forcibly to the same conditions. This may be done regularly — e.g. at the end of a television scanning line — or by sending absolute PCM values from time to time, when the nature of the signal permits. The latter is a form of *hybrid* DPCM — hybrid in the sense that both absolute and difference values are transmitted¹³. (Note however that this term, hybrid DPCM, is used by various workers to describe many very different types of systems.) Generally speaking, resetting techniques of either type are only really suitable when the predictor is fairly simple, taking only a few past sample values into account in forming the prediction.

Another way to keep the coder and decoder in step, avoiding the need for resetting, is to choose the predictor in such a way that the remote decoder output will always decay towards the (desired) locally-decoded value after a disturbance. The disturbance might be the process of switching-on or of first connecting coder and decoder together. In either case the predictors in coder and decoder will not have the same stored information of past events. Other less drastic disturbances are caused by transmission errors.




Fig. 4 - DPCM coder and decoder.

In any case, if the remotely-decoded and locally-decoded outputs are to converge to the same result then the coder and decoder predictors must be such that the influence of the stored information corresponding to a particular past instant (such as the disturbance) is gradually diluted and replaced by newer information. Resetting is the direct approach whereby exclusively PCM (absolute) information is forced into the predictor memories, flushing out the DPCM-derived values. In general, some *absolute* information must be transmitted somehow. If a predictor is 'leaky', i.e. its gain is less than unity so that it does not predict even steady d.c. conditions perfectly, then the transmitted difference values will in effect contain some d.c. — or absolute — information. This is a means of achieving our desired convergence of coder and decoder. Note that not all 'leaky' predictors give rise to a stable arrangement — see Section 4.2. In this method absolute and differential information has in effect been combined linearly as a weighted sum to form the transmitted signal. Absolute information can be inserted in other ways, perhaps non-linear: the method of van Buul¹⁴ is an example of this.

With use of 'leaky' predictors, coder and decoder may not track exactly, but differ by a small 'noise' term. This arises because of the limited precision in the arithmetic of the predictor loop and is thus implementation-dependent. In a system with infinite arithmetic precision the noise term could be made zero.

The effects resulting from any errors introduced into the coded difference signals during transmission must be considered for any practical DPCM system. One of the methods just discussed must be used to keep the decoder tracking the coder in the long term. Errors will also have short-term effects on the perceived decoded quality. These disturbances *must* be balanced with the basic coding quality in optimising a DPCM system to give the best performance overall.

2.5 Non-causal prediction

The prediction discussed so far is causal: it is based on past values of the signal. Thus, in television neither samples following the present one on the same line nor those on subsequent lines (i.e. below) may be taken into account. This is not as severe a limitation as it appears since samples from the *previous field* appear all around the sample to be predicted and can be used to supplement the restricted choice available in the present field. The development of such three-dimensional predictors is the path selected and is described in later sections.

Nevertheless a non-causal interpolative method has been proposed¹⁵. In this the sample grid is divided into two sets: a subsampled lattice whose sample values are always transmitted in absolute PCM form, and the remaining samples. For each of these remaining samples a prediction is made based only on the PCM-transmitted samples (which are thus available at both coder and decoder); it is in effect an interpolated value. For each such sample the difference between actual and interpolated values is calculated, quantised and transmitted. At the decoder these received quantised differences are added onto the interpolated values to produce the sample values. No feedback loop is involved since the prediction is based only on transmitted absolute values, not previously decoded values. Thus there is no longer any restriction imposed by causality: we may base the prediction on whichever past or future samples we choose, provided only that they are taken from the subset which is transmitted as absolute PCM. Since both absolute and differential values are transmitted this method, like the very different one described in the previous section, is sometimes also described as hybrid DPCM.

The main advantages of non-causal DPCM are ease of instrumentation (no feedback implies no critical feedback loops) and reduced vulnerability to transmission errors since there is no error propagation. An error in a 'difference' value affects only the corresponding decoded sample; an error in an 'absolute' value affects the corresponding sample and a limited number of samples derived from it.

Removal of the causality conditions does not benefit the quality of prediction very much. Although future samples can now be used as well as past ones they must all be chosen from the limited subset sent as absolute PCM. If a significant reduction of bit rate is to be achieved the proportion of PCM values (sent at full accuracy) must be small and so the basis for prediction becomes sparse.

The best application of this technique is for simply-instrumented systems of modest bit-rate reduction, especially where immunity to transmission errors is valuable.

2.6 Adaptation of predictors and/or quantisers

In the simplest form of DPCM system the predictor and quantiser will both be time-invariant. Since the source signal may have a wide variety of statistics, and the susceptibility to distortion of the destination may vary in a complex way with the signal material, it is natural to propose adaptation of the predictor and/or quantiser. For example, the predictor may in fact comprise several different predictors operating in parallel, each fed with the

common input signal and each optimised to suit a particular subset of the range of source signals. The output of one of these is then chosen as the prediction to use according to some criterion.

The possible advantage of adapting the quantiser may be illustrated as follows. Suppose that the current predicted value is large, near the top of the signal range. It follows that this prediction cannot underestimate the true sample value by very much although it could well be a significant overestimate. When the prediction is small the converse is true. So there is clearly some justification for varying the quantiser according to the size of the prediction so that the quantiser becomes progressively one-sided as the prediction approaches either extreme of the signal range. (This is sometimes known as a 'sliding quantiser'¹⁶.) Thus we can see by this example that adaptation offers potential benefits, but at the price of added complexity.

One disadvantage of adaptation is that it offers another way for the coder and decoder to get out of step. In general the adaptation can be controlled in one of two ways: either *explicitly*, where the coder sends extra control information to tell the decoder what to do, or *implicitly*, where coder and decoder make independent decisions based on information available to both of them. In either case coder and decoder could mis-track. With explicit signalling, the control information transmitted could be corrupted in transmission. With implicit signalling the information which coder and decoder nominally have in common — the transmitted difference or the predictor output — may not be identical. The transmitted difference may suffer transmission errors while the predictor output may retain memory of a previous disturbance (either at switch-on or by transmission error), as discussed in Section 2.4, or contain rounding-error noise (discussed later in Section 5.4).

Explicitly-signalled adaptation also suffers the disadvantage that it increases the bit rate by virtue of the transmission of extra control information. This tends to favour control algorithms whereby the frequency of mode changes is restricted so that less control information need be sent. This means that some of the advantage of adaptation may be lost, in that it is not possible to respond rapidly to local variation of signal statistics. There may however be a compensating benefit of reducing 'dithering' between modes.

If a regular block structure is used, one mode being selected per block, then care must be exercised otherwise it may be possible to see the join between blocks.

3. APPLICATION OF DPCM TO DIGITAL VIDEO — BACKGROUND

The application of DPCM techniques to digital video signals has received much study, both within the BBC and elsewhere.

3.1 DPCM applied to monochrome signals

Early work¹⁷ studied techniques for coding monochrome-only signals. The sampling frequency used was commonly a multiple of the line-scanning frequency, giving a sampling grid of the form shown in Fig. 5.




Fig. 5 - Line-locked sampling grid.

If we denote the predicted value of sample X by $\hat{X}$, the simplest prediction we may choose is the previous sample, i.e. $\hat{X} = A$. We shall call this 1T prediction — the predictor is in effect a one-sample-period delay. This predictor can be quite effective for modest bit-rate reduction, for example to 6 bit/sample from a PCM input having 8 bit/sample.

A better, two-dimensional (2-D) predictor makes use of the two-dimensional spatial properties of a picture scene, whereby not only is X related to adjacent previous samples on the same line (e.g. A) but it is also related to adjacent samples on the previous lines (e.g. B, C, D). The simplest 2-D predictor is $\hat{X} = (A+C)/2$, and other relatively simple weighted sums of A, B, C, and D may be devised.

Television pictures also exhibit correlation in a third dimension, that of time. If a picture-locked sampling frequency is used then there is a sample in the previous picture having the same spatial position as X. The value of this sample is an exact prediction of X provided that the picture remains stationary and noise-free. The accuracy of prediction reduces as the degree of movement increases.

A general three-dimensional (3-D) predictor may use weighted combinations of samples on many lines spread over several television fields, the constraint

being that only previously-occurring samples may be used to form the prediction.

3.2 DPCM applied to colour signals

Colour pictures may be coded into video signals in two ways: *component* or *composite*.

Component-coded signals — luminance and two colour-difference signals — can be coded by three DPCM systems working individually. It may also be possible to exploit any correlation that may exist between the component signals.

Composite-coded signals need different treatment. In the case of PAL or NTSC signals there is a subcarrier which is amplitude-modulated by the colour information. In designing a predictor we must take into account the presence of this subcarrier in addition to considering the correlation properties of the original scene.

Suppose sampling takes place at n times the subcarrier frequency, f_{sc} . In a uniform area of the picture having steady luminance and chrominance values the video signal will comprise steady subcarrier, of particular amplitude and phase, added to a d.c. level. Every n th sample of this signal will be identical, so a simple predictor is the n th-previous sample. This sample is more remote from the sample to be predicted than the simple $1T$ predictor previously mentioned and is thereby less effective for the luminance part of the signal.

The use of nT predictors has been studied in the BBC^{5,13}, both $3T$ for $3f_{sc}$ sampling and $2T$ for $2f_{sc}$. Various simple adaptive schemes were also examined as was a simple 2-D scheme¹⁸. It was clear however that substantial improvement was necessary for a system for coding within 34 Mbit/s. As mentioned in the Introduction, this system made use of blanking-removal and $2f_{sc}$ sampling. Taking these into account it was necessary to code the signal (originally 8-bit PCM) with 4.5 bit/sample.

3.3 Scope for improvement

Earlier studies of DPCM within the BBC were limited in the choice of predictors and quantisers. This was partly caused by technological limitations, but it was also natural to explore simpler systems first.

The quantisers previously used were made using complicated combinations of logic gates. Only certain quantisers could be realised readily, and a small change to the characteristic implied a new circuit design.

The greater availability of semiconductor memory components has removed this limitation. The quantiser can be treated as a look-up table stored in memory — either random-access, loaded before use, or read-only, permanently programmed with the desired curve. Random-access memory is particularly suited to the laboratory development of quantisers.

More complicated predictors have been made possible by the fall in cost and increase in speed of logic devices generally. Using a repeating, modular structure similar to that adopted for transversal filters it is possible to evaluate a weighted sum of past signal samples at real-time video speeds. With this structure the number of samples included in the prediction algorithm could in principle be increased without limit if desired.

Once restrictions on the choice of quantisers and predictors have been removed, the difficulty arises of choosing suitable ones for further study from the bewildering number of possibilities. This process has been greatly facilitated by the use of a computer in conjunction with a data bank of real-picture information.

4. DEVELOPMENT TOOLS

4.1 The approach to development

Ideally the predictor and quantiser should be jointly optimised to minimise the subjectively-assessed quality impairment introduced for all real pictures. This course is impracticable for several reasons:

- (a) we do not have a universal *objective* model by which *subjective* quality may be measured;
- (b) we cannot consider all possible pictures which have been or will be used; and
- (c) the computation involved in simultaneous optimisation would be somewhat unwieldy even if a simple objective model existed.

Many simplifications have therefore been made, for which the results finally obtained provide the justification.

The first step is to design predictors which are 'likely candidates' for further development. This is done by computation using a coarse model of quality: we assume that the best predictor has the least r.m.s. prediction error. 'Best' predictors are calculated for each of a small number of carefully chosen critical test slides for which data has been stored in the computer memory. Predictors can be tested at this stage for decoder instability. A selection of quantisers is also designed using various assumptions regarding ideal behaviour.

Performance is then checked with a wider range of critical test material using real-time flexible DPCM hardware which has been programmed with the proposed predictors and quantisers. The merits of the various proposals are informally assessed. The results may suggest changes to the predictors or quantisers. The value of any changes may be computed or may be assessed subjectively. Iteration continues until satisfactory results are obtained.

The investigation reported here concentrated on non-adaptive systems, for various reasons. One was the then-current limitations of circuit technology: whereas a general-purpose predictor having very many adjustable (but not variable on a sample-by-sample basis) sample weights was realisable, a general-purpose adaptation controller was not. To be able to implement any plausible algorithm (for the purposes of experiment) more or less implies software control, and this was not then possible either in real-time or in non-real-time simulation whose final results could be viewed. It would have been possible to build hardware to implement a *specific* algorithm but this would not permit optimisation except at great expense. Nevertheless a degree of scope for expansion was provided in the experimental hardware whereby limited adaptation of predictors could be incorporated at a later date, but this was never used, except, somewhat tangentially, for some studies of resetting strategies. Size and speed of RAM available at the time of construction also ruled out quantiser adaptation.

Another reason was more philosophical. Adaptation between predictors is only necessary if no one predictor is good enough on its own. Previously reported adaptive DPCM systems used a range of very simple predictors (each having obvious limitations) and then applied complicated switching strategies for adaptation. If, on the other hand, it were to prove the case that a single, admittedly complicated, predictor could give adequate performance then the various perils of adaptation (difficulty of optimisation and potential vulnerability to transmission errors) might be avoided. It was thus attractive to pursue this option; it did not seem to have received much previous attention and was supported by early results from the computer predictor analysis.

4.2 Software tools

The software tools developed for the project comprise several programs and a data bank containing several critical stationary pictures in $2f_{sc}$ -sampled form.

The programs are used for the following tasks:

- (a) To determine the *prediction gain* available from a given order of predictor — that is to

say a predictor with a certain number of coefficients. Prediction gain is defined for our purposes as the reduction in r.m.s. prediction error compared with a reference predictor. A suitable reference in this case is the so-called $2T$ predictor, which is simply the second-previous sample. By studying the variation of prediction gain with predictor order it is possible to decide what predictor order is likely to be worthwhile for a particular application.

- (b) To calculate the coefficients of the optimum predictor of the order specified. In fact there will be as many optimum predictors as there are separate pictures in the data base plus one derived from the aggregate data.
- (c) To apply constraints to produce practical predictor designs. The optimum predictors produced in (b) could not be realised using finite-accuracy arithmetic. The coefficients must be quantised to suit the limitations imposed by the design of the coder and decoder. In particular, the coefficients of the previous and second-previous samples were restricted to multiples of $1/8$ in the experimental equipment, and it might be desired to try further restrictions to permit economy in a production system. Other constraints which could be applied are to specify the predictor response at d.c. and subcarrier frequency.

Optimising under constraints produces better results than simply truncating the coefficients produced by (b). Indeed, step (b) is not strictly necessary but it is instructive to see how much is sacrificed by satisfying the constraints.

- (d) To design quantisers. An 'optimum' quantiser — i.e. one introducing least r.m.s. quantising error on a given picture with a given predictor — can be designed, although this is in practice much less successful in matching up to subjective assessment than the predictor optimisation. Other simple routines which generate quantisers of empirically-chosen shape are available.
- (e) To assess predictor stability. Steps (a) to (c) may well have produced a predictor with suitably high prediction gain which performs well on a coder-only assessment. However, if the decoder using this predictor is unstable it will be useless since there will always be a transmission error sooner or later! The form of the decoder is the same as in a classical control system so the standard procedures of control

theory can be applied. The method used is that of a Nyquist plot where the criterion is whether or not a certain point is encircled by a locus. This procedure is somewhat time-consuming when three-dimensional predictors are used, in which case it may be easier to try the predictors out in practice if both coder and decoder are available. The real value of the theoretical method is in confirming that predictors shown by experiment to be unstable should be so and that there is no fault in the experimental apparatus.

4.3 Hardware tools

The hardware tools developed for the project comprised a flexible coder and decoder. These were constructed on a modular basis so that predictors of increasing complexity could be investigated by adding more modules of already proved design. The same modules were also used in both coder and decoder as far as possible, with a common design of inter-module wiring.

As well as allowing a very wide choice of possible predictors the design permitted a free choice of symmetrical quantisers. Diagnostic facilities were included so that the signals at various parts of the DPCM loop could be inspected and the quantiser characteristic displayed, (see Fig. 6). A split screen was also provided so that different signals could be compared.

Initially, only a coder was constructed. As mentioned in Section 2, this is sufficient for an investigation of the inherent coding loss, but it gives no information about such things as the effect of transmission errors.

As a result of the initial experiments (see the next section) it was considered worthwhile to construct a decoder of similar design. The opportunity was also taken to include some extra facilities, both those needed to maintain picture quality and those giving added experimental facilities.

5. DEVELOPMENT WORK

The development work described here took place over a long period and went hand-in-hand with the expansion of experimental and computing facilities. In consequence, the avenues explored at any given time were often determined by what facilities were then available. In particular, neither the computer nor the experimental equipment could cope with previous-field predictors in the earlier stages of the investigation. This does not make for a logical presentation of results, so history has been somewhat distorted in what follows in the hope of achieving greater clarity.

5.1 Computed prediction gains for 'optimum' predictors

A first step in determining what can be achieved, and what degree of predictor complexity is




Fig. 6 - The flexible DPCM coder in experimental use.

worthwhile can be made using the computer to determine the 'prediction gains'* for 'optimum' predictors of particular dimensions. The program used could calculate the prediction gain of an 'optimum' predictor which uses specified sample locations, say N in all. It could then determine which sample it was most appropriate to omit in order to make an $(N-1)$ sample predictor, and compute the prediction gain for that, and so on for $(N-2)$, $(N-3)$, etc. Clearly, the optimum predictors chosen in this way depend on the detail of the sample-dropping algorithm and the initial sample configuration, so that if one starts with say $N = 50$, the results for an $N = 25$ predictor may not be the best which could be obtained with 25 samples — perhaps they should be disposed differently in space (and time). Nevertheless, this program gives a preliminary indication of what can be done. Example results from using the program are summarised in Tables 1 and 2. In both cases calculations have been performed using the 'average of all slides' data bank.

Table 1: Example results of computer optimisation of 2-D predictors.

No. of samples used in predictor, N	Prediction gain, dB (w.r.t. $2T$)
19	4.106
18	4.106
17	4.083
16	4.077
15	4.059
14	4.049
13	4.034
12	4.009
11	3.968
10	3.956
9	3.841
8	3.706
7	3.485
6	3.383
5	3.341
4	2.841
3	2.840
2	2.218

Table 1 shows how the prediction gain varies as the number N of samples included in the predictor is decreased, for the case of a 2-D predictor.

For comparison, Table 2 shows the corresponding results for 3-D predictors. Since the data bank only contained stationary slides it is necessary to find a way to make some allowance for the effects of movement. A value, representing the correlation between one field and the next is therefore supplied to the program by the user. A value of unity would

* Prediction gain is defined in Section 4.2.

apply for totally stationary pictures, and would imply that parts of the prediction taken from previous fields are just as valid as those in the present field. For real pictures containing movement a lesser value is necessary so that the optimisation routine attaches greater significance to present-field information. For the example shown in Table 2 a value of 0.85 was used. In fact subsequent experience showed that a value of 0.9 caused us to design the most useful predictors — but this should not be interpreted as an assertion that the correlation between fields in real television sequences averages 0.9.

The 3-D predictors are superior as soon as a certain minimum N is exceeded, say 5. There is also a law of diminishing returns beyond say, 20 elements.

5.2 Prediction optimisation under constraints

The results just discussed showed that clear benefits could be expected from more complicated predictors, and indeed, that quite a high degree of complexity might be worthwhile. The next step was to design some particular predictors that could be examined in practice with the experimental equipment. This meant applying constraints:

- of an instrumental nature — the experimental equipment was limited in the values of coefficients which could be used. The first two samples in the predictors, corresponding to $1T$ and $2T$ delays could only be multiplied by coefficients which were multiples of $1/8$. The remainder were less restricted, to the series $\pm 1/32, 1/16, 3/32, 1/8, 5/32, 3/16, 7/32, 1/4, 5/16, 3/8, 7/16, 1/2, 5/8, 3/4, 7/8$ and 1.
- considered appropriate to cater for certain exacting, but real situations. These impose limitations of the gain $P(f)$ of the predictor at d.c. and subcarrier frequency. It is well known that PCM coding of video signals can cause 'contouring' effects in plain areas of slowly-changing level, e.g. a sawtooth test pattern or a view of a plain wall in shadow. These effects are the main reason why 8 bit/sample are commonly used in PCM coding. Now consider a similar situation with a DPCM coder. Non-linear quantising is used where the coarseness of quantising increases as the prediction error departs from zero. Suppose that the d.c. gain of the predictor is less than unity, but the appropriate constant is added so that the prediction error is zero for plain mid-grey areas. Slowly changing areas that are, on average, very bright or very dark grey will cause significant prediction errors, the more so

Table 2: Example results of computer optimisation of 3-D predictors.

No. of samples used in predictor			Prediction gain, dB (w.r.t. 2T)
Total, N	Same-field	Previous-field	
50	19	31	6.544
49	19	30	6.544
48	19	29	6.543
47	18	29	6.542
46	18	28	6.542
45	18	27	6.541
44	18	26	6.540
43	18	25	6.538
42	18	24	6.538
41	17	24	6.534
40	16	24	6.532
39	15	24	6.528
38	14	24	6.521
37	13	24	6.520
36	13	23	6.513
35	13	22	6.513
34	13	21	6.505
33	13	20	6.504
32	13	19	6.492
31	13	18	6.475
30	13	17	6.455
29	13	16	6.445
28	13	15	6.427
27	13	14	6.409
26	12	14	6.391
25	11	14	6.385
24	11	13	6.362
23	11	12	6.343
22	11	11	6.331
21	10	11	6.294
20	9	11	6.285
19	8	11	6.219
18	8	10	6.154
17	8	9	6.035
16	8	8	5.981
15	8	7	5.945
14	7	7	5.817
13	7	6	5.785
12	6	6	5.692
11	5	6	5.692
10	5	5	5.656
9	5	4	5.605
8	5	3	5.157
7	4	3	4.756
6	3	3	4.755
5	2	3	3.810
4	2	2	2.062
3	2	1	2.046
2	2	0	1.333

as the d.c. gain is reduced, and, being thus subjected to coarser quantising, such areas will suffer contouring. The d.c. gain must therefore be kept near to unity, otherwise all the dynamic range of the quantiser is taken up by d.c. level changes in the picture, leaving less headroom for detail information. A similar argument suggests that the gain at colour subcarrier frequency should also approach unity in order to portray highly-saturated colour areas. On the other hand it is undesirable for the gain at any actual frequency to equal unity since unstable decoder behaviour will result — note that the former simple nT predictors needed frequent resetting.

In fact the situation regarding predictor gain at colour subcarrier frequency is still more complicated. As far as *stability* assessment is concerned it is the gain at colour subcarrier frequency which is significant. Let the sample to be predicted be S_0 , S_r the r th-previous sample and C_r the corresponding coefficient. The prediction is

$$\sum_{r=1}^N S_r C_r,$$

the d.c. gain is

$$\sum_{r=1}^N C_r.$$

and the gain at f_{sc} is

$$\left(\sum_{r \text{ even}} C_r - \sum_{r \text{ odd}} C_r \right).$$

That some of the samples occur on different lines and fields is irrelevant. On the other hand, the accuracy of *prediction* of plain coloured areas must take into account the lines from which the samples are taken. The effect of sub-Nyquist sampling of PAL is to make in effect a line-sequential system, with samples conveying $(U+V)$ information on one line, and $(U-V)$ on the next. Clearly the most accurate prediction for a plain coloured area will be given when:

$$\left(\sum_{r \text{ even}} C_r - \sum_{r \text{ odd}} C_r \right) \rightarrow 1$$

samples on even lines

$$\text{and } \left(\sum_{r \text{ even}} C_r - \sum_{r \text{ odd}} C_r \right) \rightarrow 0.$$

samples on odd lines

— assuming that the line containing the sample to be predicted is defined as even.

Examination of the 'optimum' predictors designed *without* applying constraints gave some interesting pointers, in that:

- (i) the d.c. gain was usually very close to 1;
- (ii) the 'even-line' colour gain was not necessarily close to 1 — typically 0.75 to 0.85 for the 'average of all slides' database, and much lower for some (not very colourful) slides;
- (iii) the 'odd-line' colour gain was typically small but non-zero;
- (iv) in consequence of (ii) and (iii), the true gain at f_{sc} was less than unity.

Practical experience with trying little- and much-constrained optimisations suggested the following general results:

- (i) Constraining the d.c. gain to 31/32 (in the hope of achieving stability) does not seem to degrade picture quality significantly. For a very low-bit-rate system there might be advantages in approaching 1 more closely, at the risk of poor stability (and also rounding-noise problems, see later) but this could not be tried with the available apparatus.
- (ii) The 'even-line' colour gain need not be constrained close to unity. Highly saturated engineering test patterns (not included in the data bank) admittedly can benefit from a slight increase in 'even-line' gain compared with the unconstrained results, but predictors giving satisfactory results over the whole range of real picture material can be found with 'even-line' gain around 0.75.
- (iii) The 'odd-line' colour gain need not be constrained to zero.

The effect of (ii) and (iii) is that relaxing the impositions made on colour behaviour permits better luminance behaviour, which was found to be more important. This was illustrated when comparing a plot of the 2-D spatial frequency prediction-error responses of two predictors. One where the colour gains were more or less left to find their own level in the optimisation had a very smooth response while another constrained to have high 'even-line' gain and zero 'odd-line' gain had a very uneven and peaky response.

5.3 Quantisers

As described in Section 2, the quantiser reduces the bit rate by limiting the accuracy with which the signal is transmitted. For video signals we try to establish that accuracy is lost only for rarely-occurring signals and in such a way that the ill-effects cannot be seen. With a particular bit rate as a target, the number of decision levels is fixed and the design problem is to dispose them in the way giving the best balance between different sorts of degradation. Roughly speaking there are three effects which it can be helpful

to distinguish — unless the number of decision levels is very small, these effects relate to different parts of the quantiser curve:

- (i) *Granular background noise* can appear in plain areas, depending on the coarseness of the quantiser in the vicinity of the origin. This, and the previously-mentioned *contouring* which can arise if the d.c. gain of the predictor is too small, behaves more or less in the manner expected from knowledge of PCM systems.
- (ii) *Edge busyness* is a term traditionally used in DPCM coding. Perhaps with more complicated predictors it would be more appropriate to talk of 'detailed-area busyness'. Areas of detail which cannot be accurately predicted appear noisy or 'busy'; prediction-error excursions are larger than in plain areas and the resulting greater coarseness of quantising causes greater quantising noise.
- (iii) *Slope overload* is again a term having origin in simpler DPCM systems. It is common for the quantiser to have a limiting characteristic — quantised differences have a maximum amplitude significantly less than the maximum signal range. Signal features that are grossly unpredictable will thus provoke severe limiting and there will be an area, usually well defined, where the quantising error is large and of predominantly one sign. Typical examples would be the edge defining a black-to-white transition. With 3-D predictors, having significant contributions from previous fields, a shot change is another example. The decoded picture following a shot change will contain gross distortions which are corrected by the feedback nature of DPCM during subsequent fields.

In addition to these effects, significant *coder* instability caused by coarse quantising has been reported¹⁹ and indeed is theoretically and practically demonstrable. However, in all of the BBC investigations no predictor/quantiser combination otherwise having any merits at all proved to be vulnerable to this. For us, therefore, it remains an interesting theoretical curiosity.

At this stage it is impossible to ignore the design of the predictor. Should one design for essentially-zero prediction error, most of the time, with very rare gross error, or for an 'all-rounder' predictor which has limited, fairly small, errors most of the time with occasional larger, but not gross errors? The latter approach proved to be more profitable, at least for a non-adaptive system. Some prediction error, with an accompanying measure of edge busyness, is tolerable in detailed areas which can mask the effect, especially

if the decision-level spacing in this intermediate region can be kept down by being able to set a fairly low limiting level.

With a given predictor, and a given set of picture data, it is possible to design an 'optimum' quantiser giving minimum r.m.s. quantising noise. Unfortunately this is not really as successful as the predictor optimisation. It cannot take proper account of masking, for example. The quantiser produced by this optimisation is roughly of the right shape of taper, but the wrong scale. Fortunately, the earlier development work concentrated on the 4½-bit quantiser needed for the 34 Mbit/s target. This has 22 levels, but because of the symmetry constraint imposed by the hardware only 11 different values had to be chosen. It was not too arduous to explore these fairly thoroughly using the hardware and the set finally chosen is listed in the Appendix. The experience gained in this work made it easier to develop a 6-bit quantiser subsequently.

5.4 Resetting, rounding error, 'smoke' and the use of rounding-error feedback

As described in Section 2, the decoder is a subset of the coder, and if we assume an interconnection which is free from transmission errors then both the 'remote' decoder and the 'local' decoder (inside the coder loop) receive the same input signals. If we suppose that at some time all the corresponding memory elements of the two decoders are forced to contain equivalent information — by some form of resetting — then the two decoders will track each other exactly, always giving the same result. This is the justification for assessing picture quality using only a coder.

When simple predictors having little memory are used it is feasible to use a resetting procedure to maintain tracking, but once the predictor is extended to three dimensions this ceases to be true. Resetting the coder and decoder predictor memories to the same value can be done either by sending a command to reset all memory instantly to a previously agreed state or by sending sufficient *absolute* information to the decoder. 'Instant resetting' is somewhat inconvenient in circuit design, but is ruled out for more fundamental reasons if the memory is large, say one field. If gross distortion is not to be apparent the common contents to which the two memories are reset must be something like the current picture information. Resetting a several-line 2-D predictor to accord with the blanking waveform* during the field blanking

* This would require the blanking interval to be present. The 34 Mbit/s system whose development is described here and in Refs. 7, 8, and 9 removed the blanking interval of the signal.

interval would work perfectly well, but there is no corresponding 'higher-order' blanking interval during which a whole field of waveform is defined.

The method of transmitting absolute information also becomes difficult. If the absolute information is sent coarsely the bit rate can remain constant — e.g. with 4 bit/sample DPCM, 4-bit PCM samples are sent instead for the duration of a reset. The distortion caused by this process need not be visible if only a few samples at a time are coarse PCM, as for the system of Ref. 13, but a whole field at a time is another matter. A brief experiment showed this to be clearly visible, especially if the resetting took place on a periodic basis. We should also note that at likely bit-error ratios there will be many errors per second so the time between resets must be short. This limitation also restricts the use of another variation in which the absolute information is sent with full accuracy. Provided this were done infrequently the average bit rate need not increase very greatly although a buffer store must be used to smooth the bit rate to match it to a continuous data channel. Another drawback revealed by brief experiment is that the sudden 'coming clean' of a picture corrupted by errors draws attention to the resetting process. Finally, the idea — related to adaptive prediction — of a hierarchy of predictors of increasing order, each being used in turn to reset the remainder was dismissed as too complex while still suffering many of the disadvantages enumerated.

The alternative to resetting, namely using 'stable' predictors, whereby the effect of any transmission error — or equally, any initial imbalance in the states of 'local' and 'remote' decoders — would decay away, therefore seemed to be the right approach. However, there is a small drawback which was revealed when both coder and decoder were completed and interconnected. When they were set to demonstrate a complicated but stable predictor the decoder did indeed recover from arbitrary switch-on conditions but its output never exactly matched that of the 'local' decoder, having a background of noisy patterning of an almost smoke-like character. If forcible resetting was applied then the patterning disappeared until the tracking of coder and decoder was disturbed by injecting a transmission error.

The reason for these effects is that the arithmetic of the predictors is performed to finite accuracy. The concept of stable predictors derives from a simple model of the decoder as a classical linear control system made of delays and multipliers. The practical decoder is slightly non-linear, since quantising error is introduced by the limited precision of the predictor arithmetic. This is distinct from the 'deliberate' quantising error introduced by the non-

linear quantiser of the coder, and to distinguish the two the term 'rounding error' will be used henceforth for this new effect due to the predictor arithmetic. The rounding error arises if any of the coefficients in the predictor has a non-integer value — which is true for almost any useful predictor excepting the nT case. Once this is true the output of the predictor requires more bits per sample to specify the output without approximation than are used at the input. Unfortunately however, the output of the predictor is fed back to the input so approximation is inevitable. If we assume that accuracy is maintained within the predictor and rounding occurs only at its output then the model as shown in Fig. 7 can be used.

In the figure the practical predictor inaccuracy is modelled by adding the error signals r, r' at coder and decoder respectively. The idealised predictor outputs are p and p' , while the rounded versions are s and s' where $s = p + r$ and $s' = p' + r'$. The idealised predictor output p is a weighted sum of past values of its input, y , which we can represent symbolically as $p = P(y)$, where P is an operator. The possibility of injecting transmission errors t is also shown.

Two things can be noted from Fig. 7. First, the only error in the *locally*-decoded output is q , the quantising error introduced by the (deliberate) quantiser Q . The influence of predictor rounding error r on the locally-decoded output is very small — it could arise since q , the quantising error introduced by Q , is dependent on its input signal $(x-s)$, which is influenced weakly by r . More importantly, after a little manipulation we can derive the *remotely*-decoded output:

quantised difference

$$\begin{aligned} &= (x - s + q) \\ &= x + q - (p + r) \\ &= x + q - r - P(x + q) \\ &= (1 - P)(x + q) - r \end{aligned}$$

received difference

$$= (1 - P)(x + q) - r + t$$

remotely-decoded signal, y'

$$\begin{aligned} &= \text{received difference} + s' \\ &= (1 - P)(x + q) - r + t + (r' + p') \\ &= (1 - P)(x + q) - r + t + r' + P(y') \end{aligned}$$

$$\text{i.e. } y' = (x + q) + \frac{(r' - r + t)}{(1 - P)}$$

If we suppose that there are no transmission errors ($t = 0$) and that resetting has taken place so




Fig. 7 - Block diagram of DPCM system, showing where rounding and transmission errors are introduced.

that r and r' are identical then $y' = (x + q)$ and the local and remote decoders give identical results. In practice this happy equilibrium will always be disturbed and we are then left with the additional error term

$$\frac{(r' - r + t)}{(1 - P)}$$

Suppose that the rounding error is white (i.e. has a flat spectrum), and for the time being assume no transmission errors. The spectrum of the error term is then governed by the factor

$$\frac{1}{1 - P(f)}$$

where $P(f)$ is now the frequency response of the predictor. We are forced to the somewhat paradoxical inference that the *better* the predictor (i.e. $P(f) \rightarrow 1$) the *worse* the effect of rounding noise ($1/(1 - P(f)) \rightarrow \infty$). If we constrain the predictor to have d.c. gain $P(0) \rightarrow 1$ then the l.f. noise increases without limit. More practically, if $P(0) = 31/32$ — a value favoured in experiment for maintaining q at an acceptable level, and also permitting the design of P to give stable response to transmission errors — then the low-frequency noise is magnified by 32. Suppose input PCM samples, and the predictor input, have 8-bit

accuracy. The output of the predictor has bits of lesser significance in addition, 5 for the case where the coefficients are multiples of $1/32$ (as in the experimental equipment). These are discarded in the rounding process, so that r has the amplitude of 'normal' 8-bit quantising noise. At low frequencies this is multiplied by 32 and is thus equivalent to 3-bit accuracy. If the two noise sources r and r' are effectively independent a further 3 dB increase applies in computing the total noise at the decoder which thus has l.f. quantising noise roughly equivalent to $2\frac{1}{2}$ bit PCM.

It was possible to test this explanation. By digitally subtracting the 'remote' and 'local' decoded outputs (with due allowance for relative instrumental delay) the rounding noise could be separated and examined. Its spectrum could be compared with that expected from knowledge of $P(f)$ and showed reasonable agreement. The predictor arithmetic could not be made more accurate with the apparatus available, but the rounding error could be increased in 1-bit increments by disabling the least-significant-bit outputs of the predictors. When this was done the expected increase of 6 dB per bit was observed until taken to the extreme where picture-dependency obviously predominated over the 'white-noise' assumption.

This suggests one way of reducing the problem, namely to improve throughout the accuracy of the predictor arithmetic. The number of bits ultimately discarded remains the same for the same coefficient resolution, but the resolution of the samples fed into the input of the predictor is increased, see Fig. 8. This method could not be applied to the experimental equipment and would in any case be difficult to apply to sufficient depth to solve the problem unaided — an extra 5 or 6 bits of input and output resolution would be difficult to achieve within the timing constraints.




Fig. 8 - Rounding error in a DPCM coder.

- (a) Simplified diagram of original arrangement, showing where accuracy is lost.
 (b) The same arrangement, with six extra bits of accuracy to render rounding noise insignificant.

NOTATION: The numbers indicate the signal representation, for example "8.6" means there are eight bits before the binary point, and six bits after it, where 1 unit is $1/256$ of the input coding range. 'S' indicates the presence of a sign bit for bipolar signals.

There is another way out, which is the use of rounding-error feedback. The discussion has assumed that r and r' are white. If their spectra could be shaped to match $1 - P(f)$ then the rounding-noise spectrum at the remote decoder output becomes white. An approximation which at least removes the

embarrassment of low-frequency noise is achieved by the rounding-error feedback arrangement of Fig. 9. The discarded low-significance bits are added back onto the predictor output after a delay. The experimental equipment was arranged so that the delay could be 1 or 2 sample periods (referred to as $1T$ or $2T$ respectively). Either had the effect of changing the low-frequency 'lumpiness' of the rounding-noise into a finer type of patterning. The exact form of the pattern obviously depended on the predictor choice, and on whether $1T$ or $2T$ rounding-error feedback was used. Generally $1T$ was preferred.

When $1T$ rounding-error feedback was used the effects of rounding noise were greatly reduced, but not sufficiently to solve the problem unaided — the difference between local- and remotely-decoded outputs was still detectable by eye, especially when the picture contained uniform areas of colour. It was therefore concluded that an increase of arithmetic resolution was desirable in addition. A simple experiment was performed with the predictor chosen for 34 Mbit/s coding. The rounding noise was separated by subtraction (as previously described), fed through an attenuator and added back onto the *analogue* video signal. In this way rounding noise having the correct spectrum (and to some extent picture dependency) could be added to the picture at a level corresponding to any desired level of processing accuracy. It was concluded that 1 extra bit of accuracy (6 dB attenuation) was not quite sufficient but that 2 bits (12 dB) would suffice. It is therefore recommended that for systems having predictors of similar characteristics the preferred arrangement is to have 10-bit input and output accuracy to the predictor, combined with $1T$ rounding-error feedback.

5.5 Systems selected for practical applications

Although many general observations were made on the way, the main thrust of the work was aimed at designing a $2f_{sc}$ DPCM system for incorporation into a 34 Mbit/s transmission system. This implied the use of a $4\frac{1}{2}$ -bit, 22-level quantiser and it was accepted that a fair degree of complexity (including a three-dimensional predictor with many coefficients) would be necessary to achieve adequate results.

Subsequently, examination of the BBC's needs in more detail prompted the BBC Communications Department to promote the concept of a 68 Mbit/s package carrying one digitally-coded video signal, its accompanying (stereo) sound plus considerable additional capacity for use for Radio programme, telephone, and other data services. Suitable capacity would result if a $2f_{sc}$ 6-bit DPCM system were used for the video signal. In this case it would be difficult




Fig. 9 - Rounding-error feedback arrangement.

to justify great complexity since a very simple 6-bit 2T system⁵ was known to offer good results — could this be improved at moderate cost? Clearly with 6 bit/sample the predictor chosen for the 4½-bit case would be better than needed so the most cost-effective simplification which retained adequate picture quality was sought. The costly items are the storage required to use the previous-field samples, the processing involved in each coefficient multiplication and addition, and, of less importance, the power supply and racking associated with these circuits. The choice is thus between

- (i) using 2-D instead of 3-D prediction,
- (ii) keeping 3-D prediction but using less samples and
- (iii) using 2-D prediction and reducing the number of samples.

The 3-D predictor chosen for the 4½-bit system uses 14 samples to form the prediction; for 6 bits with 2-D prediction a similar number of samples was still found to be necessary for adequate picture quality. On the other hand, using 3-D prediction, similar quality was achieved with only 8 samples in the prediction. Furthermore the coefficient values were, as far as possible, constrained to be easy to implement. This appeared to be a cheaper solution even at the time of the development; with the cost of memory devices falling faster than that of the devices used in processing, the advantage of this simple, 3-D predictor was clear.

Two systems have thus been developed for practical composite colour signal applications: one, for a 34 Mbit/s transmission package, using 4½ bit/sample with a 14-element, 3-D predictor and the other, for a 68 Mbit/s transmission package, using 6 bit/sample with an 8-element, 3-D predictor.

Fig. 10 illustrates the prediction algorithm for 34 Mbit/s, while the predictor coefficients and quan-

tising laws for both systems are listed in the Appendix. Subjective tests were made of both systems, and are reported in Section 7; the 68 Mbit/s proposal was subsequently taken up as a pilot scheme subjected to an extended field trial^{20,21}.

6. FURTHER WORK

The work presented here was aimed specifically at the problem of reducing the bit rate for transmission of signals which were already in composite PAL form. Indeed for distribution to terrestrial television transmitters this scenario will probably remain unchanged for some years. In contrast, the emphasis within studios is on the use of component coding of colour signals recording to CCIR Recommendation 601. One of the reasons for this is to obtain better quality in the use of special effects such as colour-separation overlay. Clearly, there will arise a need to convey signals from point to point in this format, for contribution purposes at least, and much work is being devoted to this in laboratories around the world. Many of the techniques described in this Report are equally applicable to this new problem.

Some studies have been performed, leading towards a proposal for an international standard for a contribution-quality *YUV* system operating at around 140 Mbit/s. This system uses 2-dimensional prediction and an adaptive quantising technique similar to that of Van Buul¹⁴.

7. SUBJECTIVE TESTS

7.1 The scope of the tests

Since the purpose of the investigation was to optimise the subjective quality achieved with reduced-bit-rate video coding systems, in an ideal world subjective tests would have been used throughout the investigation to guide the direction of development. To




Fig. 10 - Isometric diagram showing the sample weights of the complex 3-D predictor proposed for a 34 Mbit/s system.

do this on a formal basis is clearly impracticable, so, as described, the development work was guided by the informal assessments made by a small number of engineers associated with the project. Formal subjective tests were then conducted at the end of the investigation so that a value could be ascribed to the quality achieved using the various systems developed.

The DPCM systems described were developed for use with $2f_{sc}$ sub-Nyquist sampling and were inextricably linked with its use. This posed an extra complication in the subjective tests, which had two aims:

- (a) to judge the overall quality of the complete systems which had been developed (e.g. for 34 Mbit/s) — this would measure whether a reduced-bit-rate video coding system could be used for a particular application; and
- (b) to measure the impairments introduced by the variants of the particular coding techniques (e.g. how significant is the difference between 4½-bit and 6-bit DPCM systems) — the understanding gained from this would help direct any future development work.

To assess complete systems, as in (a), the criterion of excellence is clearly the unprocessed original analogue PAL picture. On the other hand, to assess DPCM coding *per se* we must ignore the basic impairment caused by the sub-Nyquist $2f_{sc}$ sampling process; in this case the appropriate criterion is an

8-bit PCM signal, sampled at $2f_{sc}$. This serves aim (b) as far as this Report is concerned; the detailed assessment of sub-Nyquist sampling is discussed in Ref. 9.

Various systems were tested:

- (a) the 4½-bit/sample, $2f_{sc}$ system developed for use in a 34 Mbit/s package;
- (b) the 6-bit/sample, $2f_{sc}$ system developed for use in a 68 Mbit/s package, but having a video bit rate of 53.2 Mbit/s (blanking NOT removed);
- (c) for comparison with past work, the 6-bit/sample, $2f_{sc}$ system as developed by Devereux⁵, which had the same bit rate as (b);
- (d) a 4-bit/sample, $2f_{sc}$ system (to check the effects of going beyond the 34 Mbit/s target) — note that this system had not received the same degree of serious development as (a) and (b);
- (e) a commercially-produced 68 Mbit/s system which used simple 5 bit/sample DPCM in conjunction with $3f_{sc}$ sampling.

For brevity the following brief titles will be used for the systems:

- (a) 34M;
- (b) NEW6;
- (c) OLD6;
- (d) NEW4;
- (e) SYSTEMX;

together with PCM8 for $2f_{sc}$, 8-bit PCM and ANALOG for unprocessed analogue PAL. Shortage of time, in particular the brief availability of SYSTEMX, meant that it was not possible to include all combinations in every experiment.

7.2 Subjective test method

Test pictures (both moving and stationary) were shown to several different viewing groups, each having six observers who were seated at four and six times picture height from the monitor. The viewing conditions were those of CCIR Rec. 500-1. The observers were all engineers most of whom were experienced in the critical assessment of television pictures. Several test sessions were necessary in order to accommodate all the test pictures, observers and systems.

At the beginning of each test session the observers were told the type of impairment to expect and were asked to grade picture quality. A continuous grading scale was used, based on the CCIR 5-grade quality scale. It appeared on the test forms as illustrated in Fig. 11, and was marked by the observers according to their assessments.




Fig. 11 - Appearance of grading scale on form used in subjective tests.

A 'double-stimulus' technique was used for tests. Each test session consisted of a series of presentations, each containing two test conditions, A and B, shown in the order A, B, A, B. One of the conditions in each presentation was a 'reference' condition, either unprocessed analogue PAL or 8-bit PCM sampled at $2f_{sc}$, as appropriate. Whether A or B was the reference condition was changed at random through the session, in an order unknown to the observers. After each presentation, the observers graded both conditions A and B according to the scale of Fig. 11. Each combination of particular conditions A and B with a particular picture occurred twice in the test session in random order.

The stationary picture signals were derived from a flying-spot 35 mm slide scanner, using the

following slides from the set recommended by the EBU: 'Tree,' 'Boats', 'Blackboard with toys', and also BBC Test Card F.

Two motion sequences were used: 'Shirt' and 'Trains'. These were both specially recorded using a camera and an experimental digital YUV recorder, the signals only being coded into composite PAL form on replay. This arrangement was used because the recording was also intended for use in other experiments in which the original YUV components were necessary; it produced signals of very high quality.

The 'Shirt' sequence was very critical: a patterned shirt and finely-textured tie were moved laterally in simple harmonic motion. The pattern included lines at a variety of angles and occupied most of the screen.

The momentary pauses at the extreme points of the displacement helped to reveal any additional defects introduced as a result of the motion.

The 'Trains' sequence was also fairly critical, but less so than 'Shirt'. It consisted of views of a model railway, with one moving train and other trains and a background which were stationary. It included zooming of the camera viewing angle. This sequence was much more representative of real programme material, including a variety of detail and movement directions.

Source signals, whether of stationary or moving pictures, were coded into composite PAL form using a PAL coder which incorporated a luminance notch filter.

The decision to use a notch followed tests made during the optimisation of the comb filters⁹; these indicated a clear preference, amounting to 0.5 grade, for the use of notched — over unnotched — PAL coding when viewing a colour monitor. Subsidiary tests, also using a colour monitor, have revealed, however, that the relative impairments introduced by the 34 Mbit/s system are substantially the same regardless of whether the luminance notch is used or, as is current practice, it is not. The results for viewing a monochrome monitor are more complicated and are discussed in Ref. 9.

Before discussing the detailed results, two general observations may be made. First, the quality grades given for the unprocessed PAL pictures were significantly less than Grade 5. The observers were not told which pictures were the reference, unprocessed condition and thus graded them as they thought fit. The basic impairments of even the very best PAL

pictures were thus considered quite significant by the observers, several of whom were acquainted with the quality obtained in direct *RGB* pictures. Secondly, viewers at 6H generally gave better grades than those at 4H but this applied to both original and processed pictures and thus does not affect the conclusions. Only the mean of 4H and 6H results is given for brevity. Results are listed in strict order of preference, but rounded to 0.1 grade.

7.3 Experiment 1: moving pictures, colour viewing, PCM reference

This experiment used the 'Shirt' and 'Trains' moving-picture sequences, viewed on a colour monitor and using 8-bit, $2f_{sc}$ PCM as the 'reference' condition. (Note that subjects did not know when they were viewing a 'reference', as explained in the previous section.) This test thus judges primarily the effect of different DPCM systems, of which all the newly-developed ones could be expected to exhibit some motion dependence. Twelve observers were used in all.

Results using 'Shirt' sequence:

System	Grade
PCM8	3.7
OLD6	3.6
NEW6	3.6
34M	3.3
NEW4	2.9

Results using 'Trains' sequence:

System	Grade
NEW6	3.2
PCM8	3.2
OLD6	3.2
34M	3.1
NEW4	3.1

These results contain no surprises: the lower the bit rate, the lower the quality — but the loss in quality was always small. As expected the more-or-less artificial and testing nature of the 'Shirt' sequence revealed some effects with the 4- and 4½-bit 3-D DPCM systems, but even here the drop in quality for the 4½-bit (34 Mbit/s) system was only 0.4 grade. Despite using 3-D prediction (with the possibility of impaired motion portrayal) the new 6-bit system was equivalent to the old.

7.4 Experiment 2: moving pictures, colour viewing, analogue reference

Having compared the DPCM systems in Experiment 1, this experiment used an analogue reference and aimed to assess the complete systems as proposed for 68 Mbit/s and 34 Mbit/s. Twelve observers were used.

Results using 'Shirt' sequence:

System	Grade
ANALOG	3.8
PCM8	3.8
NEW6	3.7
34M	3.4

Results using 'Trains' sequence:

System	Grade
ANALOG	3.4
PCM4	3.4
NEW6	3.4
34M	3.3

These confirm that both of the proposed systems cause little degradation. The 68 Mbit/s system is, as expected, better than the 34 Mbit/s system. Even so, the 34 Mbit/s system causes only 0.4 grade degradation with the highly critical 'Shirt' sequence.

7.5 Experiment 3: moving picture, monochrome viewing, PCM reference

This repeated Experiment 1 but viewing on a monochrome monitor. Twelve observers were used.

Results using 'Shirt' sequence:

System	Grade
PCM8	3.6
NEW6	3.6
OLD6	3.5
34M	2.9
NEW4	2.4

Results using 'Trains' sequence:

System	Grade
PCM8	3.2
34M	3.2
OLD6	3.2
NEW6	3.2
NEW4	3.1

These show that monochrome viewing was more critical when viewing the 'Shirt' sequence, but not on the more typical 'Trains' sequence.

7.6 Experiment 4: still pictures, colour viewing analogue reference

This experiment continued the assessment of complete systems started in Experiment 2, and was made more interesting by the presence of the commercial SYSTEMX. Results are given for Test Card F and also the average of all slides used. Eighteen observers were used.

Test Card F		All slides
ANALOG	3.9	3.9
NEW6	3.8	3.8
34M	3.7	3.6
SYSTEMX	3.0	3.4

These confirm that both the proposed BBC systems caused little impairment, the 68 Mbit/s proposal being marginally better than the 34 Mbit/s. The commercial 68 Mbit/s system was worse than the BBC 34 Mbit/s proposal.

7.7 Experiment 5: still pictures, monochrome viewing, analogue reference

This is a repeat of Experiment 4, but viewing in monochrome. At this stage only six observers were available to take part.

Test Card F		All slides
ANALOG	3.2	3.2
NEW6	3.2	3.1
34M	2.9	2.9
SYSTEMX	2.7	2.8

7.8 Experiment 6: effect of random transmission errors

These tests were made somewhat informally because of shortage of staff at the time of the tests, and also because of the difficulty posed by the duration required. At low bit-error ratios the time for a presentation to contain a significant number of error events is inconveniently long. The 5-point *impairment* scale was used with 6 observers present.

Two systems were tested: SYSTEMX and an approximation to the 34 Mbit/s BBC proposal. For the latter, the coded DPCM data were carried through an old experimental multiplexer/demultiplexer⁶ originally used in 60/120 Mbit/s field trials. The multiplex format was not therefore the same as might

be used in a practical 34 Mbit/s system, but the effect of errors should be substantially the same. The benefits of using a simple error corrector (Wyner-Ash (16,15)²²) were also tested as the error codec was available. In practice, it is expected that future systems will use a more powerful code such as the Reed-Solomon (63,59) code.

The results are presented in Fig. 12.

From these results we see that the 34 Mbit/s system is consistently 1 or 2 grades superior to the commercial SYSTEMX. Furthermore, the addition of a simple error corrector of only 7% redundancy ensures completely unimpaired pictures up to a b.e.r. of about 10^{-5} . It is estimated that the Reed-Solomon (63,59) code, of similar redundancy, could extend this to a b.e.r. of about 10^{-4} .




Fig. 12 - Impairment resulting from transmission errors.

— x — x — SYSTEMX
 - - Δ - - Δ - 34M with Wyner-Ash (16,15) error correction
 — ● — ● — 34M

8. CONCLUSIONS

DPCM is a useful technique for bit-rate reduction. It exploits the properties of both the source and destination of the signal to be transmitted. A DPCM coder forms a *prediction* of each sample which is based on past samples, and from it forms a prediction error which is *quantised* using a *non-linear quantiser* so that it may be transmitted using less bit/sample than the original signal.

Although worthwhile bit-rate savings for vision signals have been reported in the past using simple DPCM systems, significant improvements result when more complicated predictors are used. Use of two or three-dimensional predictors is possible, even with sub-Nyquist-sampled composite PAL video signals. The use of three dimensions (using previous-field samples

as part of the prediction) is especially valuable and permits the use of $4\frac{1}{2}$ bit/sample with very good picture quality when a 14-element predictor is used. The corresponding bit rate is low enough to be used as the vision part of a 34 Mbit/s package. For higher bit rates a less complicated predictor can be used, but 3-D prediction is still worthwhile.

The design of complicated predictors is greatly facilitated by the use of computer optimisation. A useful optimisation criterion for this is to minimise the r.m.s. prediction error. Candidate predictors produced in this way can be finally fine-tuned using subjective judgement with real-time hardware. Similar computer optimisation techniques for quantisers are less successful and a more empirical technique must be used.

A well-designed complicated predictor suffices without any need for adaptation, for the bit-rate target discussed here. It may be that adaptation of predictor or quantiser will be helpful for still lower bit rates, but various pitfalls should be noted.

It is possible to design complicated predictors which are *stable*, that is, decoder resetting is not needed and the effects of any transient disturbance die away. Whether or not a given predictor will give rise to decoder instability is mathematically determinable by the well-established Nyquist stability criterion of control theory.

Residual noise caused by the finite accuracy of predictor arithmetic can arise wherever non-unity predictor coefficients are used. It is not observed in coder-only simulation and becomes serious for predictors having d.c. gain approaching unity. The effects can be reduced to a sufficient extent by a combination of the rounding-error feedback technique and increased arithmetic accuracy in the predictor.

Two systems have been developed with practical applications in mind and were assessed by subjective tests. One is for use in a 34 Mbit/s transmission package, the other for a 68 Mbit/s one. Both offer adequate quality for distribution of PAL video signals. The 68 Mbit/s system has subsequently been assessed in a field trial.

9. ACKNOWLEDGEMENTS

The work described here took place over a long period and valuable contributions were made by many people who were members of the team. T.A. Moore of Mathematics Unit nonetheless merits particular mention for his major contributions, both to theoretical understanding and to the development of computing facilities, especially the

optimisation software.

The author also wishes to pay his personal tribute to the late G.C. Wilkinson who played a major part in constructing the equipment and conducting experiments, and whose good humour was a constant inspiration.

10. REFERENCES

1. OSBORNE, D.W., 1976. Experimental Digital Transmission of Multiplexed Video and Audio Signals at 120 Mbit/s. BBC Research Department Report No. BBC RD 1976/16.
2. REID, D.F., STOTT, J.H., and OSBORNE, D.W., 1977. Experimental Digital Transmission of Multiplexed Video and Audio Signals at 60 Mbit/s through a Satellite. BBC Research Department Report No. BBC RD 1977/32.
3. GILCHRIST, N.H.C., 1978. Digital Transmission of Multiplexed Video and Audio Signals at 140 Mbit/s over an Experimental Optical-fibre Transmission System. BBC Research Department Report No. BBC RD 1978/32.
4. WELLS, N.D., 1980. Digital Transmission Tests of Multiplexed Video and Audio Signals at 60 Mbit/s via the Orbital Test Satellite (OTS). BBC Research Department Report No. BBC RD 1980/16.
5. DEVEREUX, V.G., 1975. Digital Video: Sub-Nyquist Sampling of PAL Colour Signals. BBC Research Department Report No. BBC RD 1975/4.
6. STOTT, J.H., 1976. 60/120 Mbit/s Multiplexing Equipment for High-quality Video and Audio Signals. BBC Research Department Report No. BBC RD 1976/31.
7. RATLIFF, P.A., and STOTT, J.H., 1983. Digital Television Transmission: 34 Mbit/s PAL Investigation. BBC Research Department Report No. BBC RD 1983/9.
8. MARSDEN, R.P., 1983. Digital Video Bit-rate Reduction: Blanking Elimination for Sub-Nyquist ($2f_{sc}$) PAL. BBC Research Department Report No. BBC RD 1983/6.
9. WELLS, N.D., 1985. Digital Video: Assessment of Comb-filter Design for Sub-Nyquist PAL Codecs. BBC Research Department Report No. BBC RD 1985/18.

10. SHANNON, C.E., 1948. A Mathematical Theory of Communication. Bell Syst. Tech. J., 27, July 1948, pp 379-423.
11. HUFFMAN, D.A., 1952. A Method for the Construction of Minimum Redundancy Codes. Proc. IRE, 40, Sept. 1952, pp 1098-1101.
12. REID, D.F., 1974. Digital Video: Some Bit-rate Reduction Methods which Preserve Information in Broadcast-quality Digital Video Signals. BBC Research Department Report No. BBC RD 1974/37.
13. DEVEREUX, V.G., 1973. Digital Video: Differential Coding of PAL Signals Based on Differences between Samples One Subcarrier Period Apart. BBC Research Department Report No. BBC RD 1973/7.
14. Van BUUL, M.C.W., 1978. Hybrid DPCM, a Combination of PCM and DPCM. IEEE Trans. Comm., COM26, 3, March 1978, pp 362-368.
15. WILSON, E.J., 1981. Bit Rate Reduction to 140 Mbit/s from a Television Component Standard Requiring 216-228 Mbit/s. IBA E.D. Dept. Interim Report No. 111/81, February 1981.
16. SABATIER, J., AGUILLE, G., JOLIVET, J-C., LE BLOCH, H., and RIOU, A. 1982. Coding System for 34 Mbit/s. Digital Transmission of Colour Television. Commutation et Transmission No. 1, 1982, pp 69-82.
17. O'NEAL, J.B. 1965. Predictive Quantizing Systems (Differential Pulse Code Modulation) for the Transmission of Television Signals. Bell. Syst. Tech. J., 46, May/June 1966, pp 689-721.
18. DEVEREUX, V.G., 1975. Digital Video: Differential Coding of PAL Colour Signals using Same-line and Two-dimensional Prediction. BBC Research Department Report No. BBC RD 1975/20.
19. PIRSCH, P. 1982. Stability Conditions for DPCM Coders. IEEE Trans. Comm., COM26, 3, March 1978, pp 362-368.
20. STOTT, J.H., 1984. Digital Transmission: 68 Mbit/s PAL Field Trial System. BBC Research Department Report No. BBC RD 1984/5.
21. LEWIS, A.R., OSBORNE, D.W., and KEARSEY, B.N., 1984. Transmission of Broadcast Services over a 140 Mbit/s Digital Bearer: a Joint BBC/BT Pilot Scheme. Radio and Electronic Engineer, Vol. 54, No. 11/12, Nov/Dec 1984, pp 445-456.
22. STOTT, J.H., OLIPHANT, A., and OSBORNE, D.W., 1974. Digital Video: Error-correcting Codes and a Practical Study of a Wyner-Ash Error Corrector. BBC Research Department Report No. BBC RD 1974/40.

APPENDIX
Specifications of practical systems
 (note that these are the subject of a patent application).

A.1 34 Mbit/s system

(a) Predictor

The predictor forms a weighted sum of fourteen previous samples, six taken from the present field and eight taken from the previous field, as illustrated in Fig. 10.

The weighting coefficients C_1 to C_{14} have the following values:

Present field		Previous field	
C_1	$1/8$	C_7	$3/8$
C_2	$1/2$	C_8	$1/16$
C_3	$1/4$	C_9	$-1/4$
C_4	$1/16$	C_{10}	$3/16$
C_5	$-1/8$	C_{11}	$3/16$
C_6	$5/32$	C_{12}	$-5/32$
		C_{13}	$-3/32$
		C_{14}	$-5/16$

(b) Quantiser

The quantiser has 22 levels and is symmetrical so that the signal is quantised according to its magnitude as shown below:

Prediction-error magnitude	Quantised output
0 or 1	1*
2 or 3	2
4 or 5	4
6 to 8	7
9 to 12	10
13 to 16	14
17 to 21	19
22 to 27	24
28 to 35	31
36 to 47	40
48 to 255	54

* Prediction error of 0 or +1 is quantised as +1, -1 as -1.




Fig. A.1 - Position of samples used in DPCM predictor proposed for a 68 Mbit/s system.

A.2 68 Mbit/s system

(a) Predictor

The predictor forms a weighted sum of eight previous samples as shown in Fig. A.1.

The values of the weighting coefficients are given in the following table:

Coefficient	Value
C_1	$1/8$
C_2	$1/2$
C_3	$5/16$
C_4	$3/32$
C_5	$-3/16$
C_6	$3/8$
C_7	$1/4$
C_8	$-1/2$

(b) Quantiser

The quantiser operates on the bipolar prediction-error signal which is generated by the subtractor. The quantiser is symmetrical so that the signal is quantised according to its magnitude irrespective of sign, as specified in the following table:

Prediction-error magnitude	Quantised output
0	0
1	1
2	2
3	3
4	4
5	5
6	6
7	7
8	8
9	9
10	10
11	11
12	12
13 or 14	13
15 or 16	15
17 or 18	17
19 or 20	19
21 or 22	21
23 or 24	23
25 or 26	25
27 to 29	28
30 to 32	31
33 to 35	34
36 to 38	37
39 to 42	40
43 to 46	44
47 to 50	48
51 to 56	53
57 to 63	59
64 to 71	67
72 to 83	76
84 to 255	90

